arXiv:1605.08222v2 [cs.DC] 4 Oct 2016

ICE: A General and Validated Energy Complexity Model for Multithreaded Algorithms (IFI-UiT Technical Report 2016-77)

Vi Ngoc-Nha Tran and Phuong Hoai Ha

The Arctic Green Computing Group
Department of Computer Science
UiT The Arctic University of Norway
Tromso, Norway
{vi.tran, phuong.hoai.ha}@uit.no

September 23, 2016

Abstract

Like time complexity models that have significantly contributed to the analysis and development of fast algorithms, energy complexity models for parallel algorithms are desired as crucial means to develop energy efficient algorithms for ubiquitous multicore platforms. Ideal energy complexity models should be validated on real multicore platforms and applicable to a wide range of parallel algorithms. However, existing energy complexity models for parallel algorithms are either theoretical without model validation or algorithm-specific without ability to analyze energy complexity for a wide-range of parallel algorithms.

This paper presents a new general validated energy complexity model for parallel (multithreaded) algorithms. The new model abstracts away possible multicore platforms by their static and dynamic energy of computational operations and data access, and derives the energy complexity of a given algorithm from its *work*, *span* and *I/O* complexity. The new model is validated by different sparse matrix vector multiplication (SpMV) algorithms and dense matrix multiplication (matmul) algorithms running on high performance computing (HPC) platforms (e.g., Intel Xeon and Xeon Phi). The new energy complexity model is able to characterize and compare the energy consumption of SpMV and matmul kernels according to three aspects: different algorithms, different input matrix types and different platforms. The prediction of the new model regarding which algorithm consumes more energy with different inputs on different platforms, is confirmed by the experimental results. In order to improve the usability and accuracy of the new model for a wide range of platforms,

the platform parameters of ICE model are provided for eleven platforms including HPC, accelerator and embedded platforms.

1 Introduction

Understanding the energy complexity of algorithms is crucial important to improve the energy efficiency of algorithms [31, 30, 29, 20] and reduce the energy consumption of computing systems [28, 27, 21]. One of the main approaches to understand the energy complexity of algorithms is to devise energy models.

Significant efforts have been devoted to developing power and energy models in literature [2, 8, 7, 17, 18, 16, 24, 26]. However, there are no analytic models for multithreaded algorithms that are both applicable to a wide range of algorithms and comprehensively validated yet (cf. Table 1). The existing *parallel* energy models are either theoretical studies without validation or only applicable for specific algorithms. Modeling energy consumption of *parallel* algorithms is difficult since the energy models must take into account the complexity of both parallel algorithms and parallel platforms. The algorithm complexity results from parallel computation, concurrent memory accesses and inter-process communication. The platform complexity results from multicore architectures with deep memory hierarchy.

The existing models and their classification are summarized in Table 1. To the best of our knowledge, the proposed ICE (Ideal Cache Energy) complexity model is the first energy model that covers all three aspects: i) ability to analyze the energy complexity of parallel algorithms (i.e. Energy complexity analysis for parallel algorithms), ii) applicability to a wide range of algorithms (i.e., Algorithm generality), and iii) model validation (i.e., Validation). Section 7 describes how the ICE model complements the other currently used models.

The energy complexity model ICE proposed in this study is for general multithreaded algorithms and validated on three aspects: different algorithms for a given problem, different input types and different platforms. The proposed model is an analytic model which characterizes both algorithms (e.g., representing algorithms by their *work*, *span* and *I/O* complexity) and platforms (e.g., representing platforms by their static and dynamic energy of memory accesses and computational operations). By considering *work*, *span* and *I/O* complexity, the new ICE model is applicable to any multithreaded algorithms.

The new ICE model is designed for analyzing the energy *complexity* of algorithms and therefore the model does not provide the estimation of absolute energy consumption. The goal of the ICE model is to answer energy complexity question: *"Given two parallel algorithms A and B for a given problem, which algorithm consumes less energy analytically?"*. Hence, the details of underlying systems (e.g., runtime and architectures) are abstracted away to keep ICE model simple and suitable for complexity analysis. O-notation represents an *asymptotic upper-bound* on energy complexity.

In this work, the following contributions have been made.

- Devising a new general energy model ICE for analyzing the energy com-

Table 1: Energy Model Summary

Study	**Energy complexity analysis for parallel algorithms**	**Algorithm generality**	**Validation**
LEO [24]	No	General	Yes
POET [16]	No	General	Yes
Koala [26]	No	General	Yes
Roofline [8, 7]	No	General	Yes
Energy scalability [17, 18]	Yes	General	No
Sequential energy complexity [25]	No	General	Yes
Alonso et al. [2]	Yes	Algorithm-specific	Yes
Malossi et al. [23]	Yes	Algorithm-specific	Yes
ICE model (this study)	**Yes**	**General**	**Yes**

To the best of our knowledge, the ICE model is the first *validated model that supports* energy complexity analysis for *general multi-threaded algorithms.*

plexity of a wide range of multithreaded algorithms based on their *work*, *span* and *I/O* complexity (cf. Section 3). The new ICE model abstracts away possible *multicore platforms* by their static and dynamic energy of computational operations and memory access. The new ICE model complements previous energy models such as energy roofline models [8, 7] that abstract away possible *algorithms* to analyze the energy consumption of different multicore platforms.

- Conducting two case studies (i.e., SpMV and matmul) to demonstrate how to apply the ICE model to find energy complexity of parallel algorithms. The selected parallel algorithms for SpMV are three algorithms: Compressed Sparse Column(CSC), Compressed Sparse Block(CSB) and Compressed Sparse Row(CSR)(cf. Section 4). The selected parallel algorithms for matmul are two algorithms: a basic matmul algorithm and a cache-oblivious algorithm (cf. Section 5).

- Validating the ICE energy complexity model with both data-intensive (i.e., SpMV) and computation-intensive (i.e., matmul) algorithms according to three aspects: different algorithms, different input types and different platforms. The results show the precise prediction on which validated SpMV algorithm (i.e., CSB or CSC) consumes more energy when using different matrix input types from Florida matrix collection [10] (cf. Section 6.5). The results also show the precise prediction on which validated matmul algorithm (i.e., basic or cache-oblivious) consumes more energy (cf. Section 6.6). The model platform-related parameters for 11 platforms, including x86, ARM and GPU, are provided to facilitate the deployment of the ICE model.

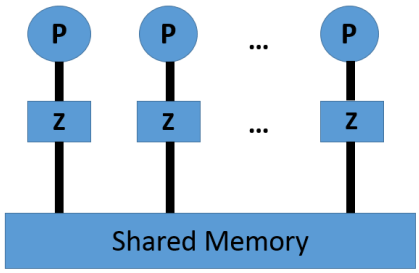

Figure 1: A Shared Memory Machine Model with Private Caches

2 ICE Shared Memory Machine Model

Generally speaking, the energy consumption of a parallel algorithm is the sum of i) static energy (or leakage) E_{static}, ii) dynamic energy of computation E_{comp} and iii) dynamic energy of memory accesses E_{mem}. The static energy E_{static} is proportional to the execution time of the algorithm while the dynamic energy of computation and the dynamic energy of memory accesses are proportional to the number of computational operations and the number of memory accesses of the algorithm, respectively [18]. As a result, in the new ICE complexity model, the energy complexity of a multithreaded algorithm is analyzed based on its *span complexity* [9] (for the static energy), *work complexity* [9] (for the dynamic energy of computation) and *I/O complexity* (for the dynamic energy of memory accesses) (cf. Section 3). This section describes shared-memory machine models supporting I/O complexity analysis for parallel algorithms.

The first memory model we consider is parallel external memory (PEM) model [3], an extension of the Parallel Random Access Machine (PRAM) model that includes a two-level memory hierarchy. In the PEM model, there are n cores (or processors) each of which has its own *private* cache of size Z (in bytes) and shares the main memory with the other cores (cf. Figure 1). When n cores access n distinct blocks from the shared memory *simultaneously*, the I/O complexity in the PEM model is $O(1)$ instead of $O(n)$. Although the PEM model is appropriate for analyzing the I/O complexity of parallel algorithms in terms of time performance [3], we have found that the PEM model is not appropriate for analyzing parallel algorithms in terms of the dynamic energy of memory accesses. In fact, even when the n cores can access data from the main memory simultaneously, the *dynamic* energy consumption of the access is proportional to the number n of accessing cores (because of the load-store unit activated within each accessing core and the energy compositionality of parallel computations [15, 22]), rather than a constant as implied by the PEM model.

As a result, we consider the ideal distributed cache (IDC) model [13] to analyze I/O complexity of multithreaded algorithms in terms of dynamic energy consumption. Since the cache complexity of m misses is $O(m)$ regardless of

whether or not the cache misses are incurred simultaneously by the cores, the IDC model reflects the aforementioned dynamic energy consumption of memory accesses by the cores.

However, the IDC model is mainly designed for analyzing the cache complexity of divide-and-conquer algorithms, making it difficult to apply to general multi-threaded algorithms targeted by our new ICE model. Constraining the new ICE model to the IDC model would limit the applicability of the ICE model to a wide range of multithreaded algorithms.

In order to make our new ICE complexity model applicable to a wide range of multithreaded algorithms, we show that the cache complexity analysis using the traditional (sequential) ideal cache (IC) model [12] can be used to find an upper bound on the cache complexity of the same algorithm using the IDC model (cf. Lemma 2.1). As the sequential execution of multithreaded algorithms is a valid execution regardless of whether they are divide-or-conquer algorithms, the ability to analyze the cache complexity of multithreaded algorithms via their sequential execution in the ICE complexity model improves the usability of the ICE model.

Let $Q_1(Alg, B, Z)$ and $Q_P(Alg, B, Z)$ be the cache complexity of a parallel algorithm Alg analyzed in the (uniprocessor) ideal cache (IC) model [12] with block size B and cache size Z (i.e, running Alg with a single core) and the cache complexity analyzed in the (multicore) IDC model with P cores each of which has a private cache of size Z and block size B, respectively. We have the following lemma:

Lemma 2.1. *The cache complexity $Q_P(Alg, B, Z)$ of a parallel algorithm Alg analyzed in the ideal distributed cache (IDC) model with P cores is bounded from above by the product of P and the cache complexity $Q_1(Alg, B, Z)$ of the same algorithm analyzed in the ideal cache (IC) model. Namely,*

$$Q_P(Alg, B, Z) \leq P * Q_1(Alg, B, Z) \tag{1}$$

Proof. (Sketch) Let $Q_P^i(Alg, B, Z)$ be the number of cache misses incurred by core i during the parallel execution of algorithm Alg in the IDC model. Because caches do not interfere with each other in the IDC model, the number of cache misses incurred by core i when executing algorithm Alg in parallel by P cores is not greater than the number of cache misses incurred by core i when executing the whole algorithm Alg only by core i. That is,

$$Q_P^i(Alg, B, Z) \leq Q_1(Alg, B, Z) \tag{2}$$

or

$$\sum_{i=1}^{P} Q_P^i(Alg, B, Z) \leq P * Q_1(Alg, B, Z) \tag{3}$$

On the other hand, since the number of cache misses incurred by algorithm Alg when it is executed by P cores in the IDC model is the sum of the numbers

of cache misses incurred by each core during the Alg execution, we have

$$Q_P(Alg, B, Z) = \sum_{i=1}^{P} Q_P^i(Alg, B, Z) \tag{4}$$

From Equations 3 and 4, we have

$$Q_P(Alg, B, Z) \leq P * Q_1(Alg, B, Z) \tag{5}$$

□

We also make the following assumptions regarding platforms.

- Algorithms are executed with the best configuration (e.g., maximum number of cores, maximum frequency) following the race-to-halt strategy.
- The I/O parallelism is bounded from above by the computation parallelism. Namely, each core can issue a memory request only if its previous memory requests have been served. Therefore, the work and span (i.e., critical path) of an algorithm represent the parallelism for both I/O and computation [9].

3 Energy Complexity in ICE model

This section describes two energy complexity models, a platform-supporting energy complexity model considering both platform and algorithm characteristics and platform-independent energy complexity model considering only algorithm characteristics. The platform-supporting model is used when platform parameters in the model are available while platform-independent model analyses energy complexity of algorithms without considering platform characteristics.

3.1 Platform-supporting Energy Complexity Model

This section describes a methodology to find energy complexity of algorithms. The energy complexity model considers three groups of parameters: machine-dependent, algorithm-dependent and input-dependent parameters. The reason to consider all three parameter-categories is that only operational intensity [32] is insufficient to capture the characteristics of algorithms. Two algorithms with the same values of operational intensity might consume different levels of energy. The reasons are their differences in data accessing patterns leading to performance scalability gap among them. For example, although the sequential version and parallel version of an algorithm may have the same operational intensity, they may have different energy consumption since the parallel version would have less static energy consumption because of shorter execution time.

The energy consumption of a parallel algorithm is the sum of i) static energy (or leakage) E_{static}, ii) dynamic energy of computation E_{comp} and iii) dynamic energy of memory accesses E_{mem}: $E = E_{static} + E_{comp} + E_{mem}$ [8, 17, 18]. The

Table 2: ICE Model Parameter Description

Machine	Description
ϵ_{op}	dynamic energy of one operation (average)
$\epsilon_{I/O}$	dynamic energy of a random memory access (1 core)
π_{op}	static energy when performing one operation
$\pi_{I/O}$	static energy of a random memory access
Algorithm	**Description**
$Work$	Number of work in flops of the algorithm [9]
$Span$	The critical path of the algorithm [9]
I/O	Number of cache line transfer of the algorithm [9]

static energy E_{static} is the product of the execution time of the algorithm and the static power of the whole platform. The dynamic energy of computation and the dynamic energy of memory accesses are proportional to the number of computational operations $Work$ and the number of memory accesses I/O, respectively. Pipelining technique in modern architectures enables overlapping computation with memory accesses [15]. Since computation time and memory-access time can be overlapped, the execution time of the algorithm is assumed to be the maximum of computation time and memory-access time [8]. Therefore, the energy consumption of algorithms is computed by Equation 6, where the values of ICE parameters, including ϵ_{op}, $\epsilon_{I/O}$, π_{op}, and $\pi_{I/O}$ are described in Table 2 and computed by the Equation 7, 8, 9, and 10, respectively.

$$E = P^{sta} \times max(T^{comp}, T^{mem}) + \epsilon_{op} \times Work + \epsilon_{I/O} \times I/O \tag{6}$$

$$\epsilon_{op} = P^{op} \times \frac{F}{Freq} \tag{7}$$

$$\epsilon_{I/O} = P^{I/O} \times \frac{M}{Freq} \tag{8}$$

$$\pi_{op} = P^{sta} \times \frac{F}{Freq} \tag{9}$$

$$\pi_{I/O} = P^{sta} \times \frac{M}{Freq} \tag{10}$$

The dynamic energy of one operation by one core ϵ_{op} is the product of the consumed power of one operation by one active core P^{op} and the time to perform one operation. Equation 7 shows how ϵ_{op} relates to frequency $Freq$ and the number of cycles per operation F. Similarly, the dynamic energy of a random access by one core $\epsilon_{I/O}$ is the product of the consumed power by one active core

Table 3: Platform parameter summary. The parameters of the first nine platforms are derived from [7] and the parameters of the two new platforms are found in this study.

Platform	Processor	ϵ_{op}(nJ)	π_{op}(nJ)	$\epsilon_{I/O}$(nJ)	$\pi_{I/O}$(nJ)
Nehalem i7-950	Intel i7-950	0.670	2.455	50.88	408.80
Ivy Bridge i3-3217U	Intel i3-3217U	0.024	0.591	26.75	58.99
Bobcat CPU	AMD E2-1800	0.199	3.980	27.84	387.47
Fermi GTX 580	NVIDIA GF100	0.213	0.622	32.83	45.66
Kepler GTX 680	NVIDIA GK104	0.263	0.452	27.97	26.90
Kepler GTX Titan	NVIDIA GK110	0.094	0.077	17.09	32.94
XeonPhi KNC	Intel 5110P	0.012	0.178	8.70	63.65
Cortex-A9	TI OMAP 4460	0.302	1.152	51.84	174.00
Arndale Cortex-A15	Samsung Exynos 5	0.275	1.385	24.70	89.34
Xeon	2xIntel E5-2650l v3	0.263	0.108	8.86	23.29
Xeon-Phi	Intel 31S1P	0.006	0.078	25.02	64.40

performing one I/O (i.e., cache-line transfer) $P^{I/O}$ and the time to perform one cache line transfer computed as $M/Freq$, where M is the number of cycles per cache line transfer (cf. Equation 8). The static energy of operations π_{op} is the product of the whole platform static power P^{sta} and time per operation. The static energy of one I/O $\pi_{I/O}$ is the product of the whole platform static power and time per I/O, shown by Equation 9 and 10.

In order to compute *work*, *span* and *I/O* complexity of the algorithms, the input parameters also need to be considered. For example, SpMV algorithms consider input parameters listed in Table 4. Cache size is captured in the ICE model by the *I/O complexity* of the algorithm. Note that in the ICE machine model (Section 2), cache size Z is a constant and may disappear in the *I/O complexity* (e.g., O-notation).

The details of how to obtain the ICE parameters of recent platforms are discussed in Section 6.1. The actual values of ICE platform parameters for 11 recent platforms are presented in Table 3.

The computation time of parallel algorithms is proportional to the span complexity of the algorithm, which is $T^{comp} = \frac{Span \times F}{Freq}$ where $Freq$ is the processor frequency, and F is the number of cycles per operation. The memory-access time of parallel algorithms in the ICE model is proportional to the I/O complexity of the algorithm divided by its I/O parallelism, which is $T^{mem} = \frac{I/O}{I/O-parallelism} \times \frac{M}{Freq}$. As I/O parallelism, which is the average number of I/O ports that the algorithm can utilize per step along the span, is bounded by the computation parallelism $\frac{Work}{Span}$, namely the average number of cores that the algorithm can utilize per step along the span (cf. Section 2), the memory-access time T^{mem} becomes: $T^{mem} = \frac{I/O \times Span \times M}{Work \times Freq}$ where M is the number of

cycles per cache line transfer. If an algorithm has T^{comp} greater than T^{mem}, the algorithm is a CPU-bound algorithm. Otherwise, it is a memory-bound algorithm.

3.1.1 CPU-bound Algorithms

If an algorithm has computation time T^{comp} longer than data-accessing time T^{mem} (i.e., CPU-bound algorithms), the ICE energy complexity model becomes Equation 11 which is simplified as Equation 12.

$$E = P^{sta} \times \frac{Span \times F}{Freq} + \epsilon_{op} \times Work + \epsilon_{I/O} \times I/O \tag{11}$$

or

$$E = \pi_{op} \times Span + \epsilon_{op} \times Work + \epsilon_{I/O} \times I/O \tag{12}$$

3.1.2 Memory-bound Algorithms

If an algorithm has data-accessing time longer than computation time (i.e., memory-bound algorithms): $T^{mem} \geq T^{comp}$, energy complexity becomes Equation 13 which is simplified as Equation 14.

$$E = P^{sta} \times \frac{I/O \times Span \times M}{Work \times Freq} + \epsilon_{op} \times Work + \epsilon_{I/O} \times I/O \tag{13}$$

or

$$E = \pi_{I/O} \times \frac{I/O \times Span}{Work} + \epsilon_{op} \times Work + \epsilon_{I/O} \times I/O \tag{14}$$

3.2 Platform-independent Energy Complexity Model

This section describes the energy complexity model that is platform-independent and considers only algorithm characteristics. This complexity model is used when analyzing energy complexity of an algorithm without platform parameters. When the platform parameters (i.e., ϵ_{op}, $\epsilon_{I/O}$, π_{op}, and $\pi_{I/O}$) are unavailable, the energy complexity model is derived from Equation 6 because the platform parameters are constants and can be removed. Assuming $\pi_{max} = max(\pi_{op}, \pi_{I/O})$, after removing platform parameters, the platform-independent energy complexity model are shown in Equation 15.

$$E = O(Work + I/O + max(Span, \frac{I/O \times Span}{Work})) \tag{15}$$

4 A Case Study of Sparse Matrix Multiplication

SpMV is one of the most common application kernels in Berkeley dwarf list [4]. It computes a vector result y by multiplying a sparse matrix A with a dense vector x: $y = Ax$. SpMV is a data-intensive kernel and has irregular memory-access patterns. The data access patterns for SpMV is defined by

Table 4: SpMV Input Parameter Description

SpMV Input	Description
n	Number of rows
nz	Number of nonzero elements
nr	Maximum number of nonzero in a row
nc	Maximum number of nonzero in a column
β	Size of a block

its sparse matrix format and matrix input types. There are several sparse matrix formats and SpMV algorithms in literature. To name a few, they are Coordinate Format (COO), Compressed Sparse Column (CSC), Compressed Sparse Row (CSR), Compressed Sparse Block (CSB), Recursive Sparse Block (RSB), Block Compressed Sparse Row (BCSR) and so on. Three popular SpMV algorithms, namely CSC, CSB and CSR are chosen to validate the proposed energy complexity model. They have different data-accessing patterns leading to different values of I/O, work and span complexity. Since SpMV is a memory-bound application kernel, Equation 14 is applied. The input matrices of SpMV have different parameters listed in Table 4.

4.1 Compressed Sparse Row

CSR is a standard storage format for sparse matrices which reduces the storage of matrix compared to the tuple representation [19]. This format enables row-wise compression of A with size $n \times n$ (or $n \times m$) to store only the non-zero nz elements. Let nz be the number of non-zero elements in matrix A. The *work* complexity of CSR SpMV is $\Theta(nz)$ where $nz >= n$ and *span* complexity is $O(nr + \log n)$ [6], where nr is the maximum number of non-zero elements in a row. The *I/O* complexity of CSR in the sequential I/O model of row-major layout is $O(nz)$ [5] namely, scanning all non-zero elements of matrix A costs $O(\frac{nz}{B})$ I/Os with B is the cache block size. However, randomly accessing vector x causes the total of $O(nz)$ I/Os. Applying the proposed model on CSR SpMV, their total energy complexity are computed as Equation 16.

$$E_{CSR} = O(\epsilon_{op} \times nz + \epsilon_{I/O} \times nz + \pi_{I/O} \times (nr + \log n)) \tag{16}$$

4.2 Compressed Sparse Column

CSC is the similar storage format for sparse matrices as CSR. However, it compresses the sparse matrix in column-wise manner to store the non-zero elements. The *work* complexity of CSC SpMV is $\Theta(nz)$ where $nz >= n$ and *span* complexity is $O(nc + \log n)$, where nc is the maximum number of non-zero elements in

$$E_{CSB} = O(\epsilon_{op} \times (\frac{n^2}{\beta^2} + nz) + \epsilon_{I/O} \times (\frac{n^2}{\beta^2} + \frac{nz}{B}) + \pi_{I/O} \times \frac{(\frac{n^2}{\beta^2} + \frac{nz}{B}) \times (\beta \times \log \frac{n}{\beta} + \frac{n}{\beta})}{(\frac{n^2}{\beta^2} + nz)}) \tag{18}$$

---

a column. The *I/O* complexity of CSC in the sequential I/O model of column-major layout is $O(nz)$ [5]. Similar to CSR, scanning all non-zero elements of matrix A in CSC format costs $O(\frac{nz}{B})$ I/Os. However, randomly updating vector y causing the bottle neck with total of $O(nz)$ I/Os. Applying the proposed model on CSC SpMV, their total energy complexity are computed as Equation 17.

$$E_{CSC} = O(\epsilon_{op} \times nz + \epsilon_{I/O} \times nz + \pi_{I/O} \times (nc + \log n)) \tag{17}$$

4.3 Compressed Sparse Block

Given a sparse matrix A, while CSR has good performance on SpMV $y = Ax$, CSC has good performance on transpose sparse matrix vector multiplication $y = A^T \times x$, Compressed sparse blocks (CSB) format is efficient for computing either Ax or $A^T x$. CSB is another storage format for representing sparse matrices by dividing the matrix A and vector x, y to blocks. A block-row contains multiple chunks, each chunks contains consecutive blocks and non-zero elements of each block are stored in Z-Morton-ordered [6]. From Beluc et al. [6], CSB SpMV computing a matrix with nz non-zero elements, size $n \times n$ and divided by block size $\beta \times \beta$ has span complexity $O(\beta \times \log \frac{n}{\beta} + \frac{n}{\beta})$ and *work* complexity as $\Theta(\frac{n^2}{\beta^2} + nz)$.

I/O complexity for CSB SpMV is not available in the literature. We do the analysis of CSB manually by following the master method [9]. The *I/O* complexity is analyzed for the algorithm CSB_SpMV(A,x,y) from Beluc et al. [6]. The I/O complexity of CSB is similar to *work* complexity of CSB $O(\frac{n^2}{\beta^2} + nz)$, only that non-zero accesses in a block is divided by B: $O(\frac{n^2}{\beta^2} + \frac{nz}{B})$, where B is cache block size. The reason is that non-zero elements in a block are stored in Z-Morton order which only requires $\frac{nz}{B}$ I/Os. The energy complexity of CSB SPMV is shown in Equation 18.

From the complexity analysis of SpMV algorithms using different layouts, the complexity of CSR-SpMV, CSC-SpMV and CSB-SpMV are summarized in Table 5.

5 A Case Study of Dense Matrix Multiplication

Besides SpMV, we also apply the ICE model to dense matrix multiplication (matmul). Unlike SpMV, a data-intensive kernel, matmul is a computation-intensive kernel used in high performance computing. It computes output ma-

Table 5: SpMV Complexity Analysis

Complexity	CSC-SpMV	CSB-SpMV	CSR-SpMV
Work	$\Theta(nz)$ [6]	$\Theta(\frac{n^2}{\beta^2} + nz)$ [6]	$\Theta(nz)$ [6]
I/O	$O(nz)$ [5]	$O(\frac{n^2}{\beta^2} + \frac{nz}{B})$ [this study]	$O(nz)$ [5]
Span	$O(nc + \log n)$ [6]	$O(\beta \times \log \frac{n}{\beta} + \frac{n}{\beta})$ [6]	$O(nr + \log n)$ [6]

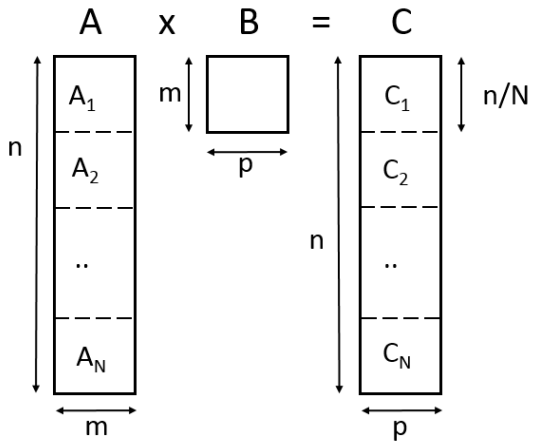

Figure 2: Partition approach for parallel matmul algorithms. Each sub-matrix A_i has size $\frac{n}{N} \times m$ and each sub-matrix C_i has size $\frac{n}{N} \times p$.

trix C (size n x p) by multiplying two dense matrices A (size n x m) and B (size m x p): $C = A \times B$. In this work, we implemented two matmul algorithms (i.e., a basic algorithm and a cache-oblivious algorithm [12]) and apply the ICE analysis to find their energy complexity. Both algorithms partition matrix A and C equally to N sub-matrices (e.g., A_i with i=(1,2,..,N)), where N is the number of cores in the platform. The partition approach is shown in Figure 2. Each core computes a sub-matrix C_i: $C_i = A_i \times B$. Since matmul is a computation-bound application kernel, Equation 12 is applied.

5.1 Basic Matmul Algorithm

The basic matmul algorithm is described in Figure 3. Its work complexity is $\Theta(2nmp)$ [33] and span complexity is $\Theta(\frac{2nmp}{N})$ because the computational work is divided equally to N cores due to matrix partition approach. When matrix size of matrix B is bigger than platform cache size, the basic algorithm loads

```
1  for i = 1 to n
2      for j = 1 to p
3          for k = 1 to m
4              C(i,j) = C(i,j) + A(i,k) * B(k,j)
```

Figure 3: Basic matmul algorithm, where sizes of matrix A, B, C are nxm, mxp, nxp, respectively.

$$E_{CO} = O(\epsilon_{op} \times 2nmp + \epsilon_{I/O} \times (n+m+p+\frac{nm+mp+np}{B}+\frac{nmp}{B\sqrt[2]{Z}}) + \pi_{op} \times \frac{2nmp}{N}) \tag{20}$$

---

matrix B n times (i.e., once for computing each row of C), results in $\frac{nmp}{B}$ cache block transfer, where B is cache block size. In total, I/O complexity of the basic matmul algorithm is $\Theta(\frac{nm+nmp+np}{B})$. Applying the ICE model on this algorithm, the total energy complexity is computed as Equation 19.

$$E_{basic} = O(\epsilon_{op} \times 2nmp + \epsilon_{I/O} \times \frac{nm+nmp+np}{B} + \pi_{op} \times \frac{2nmp}{N}) \tag{19}$$

5.2 Cache-oblivious Matmul Algorithm

The cache-oblivious matmul (CO-matmul) algorithm [12] is a divide-and-conquer algorithm. It has work complexity the same as the basic matmul algorithm $\Theta(2nmp)$. Its span complexity is also $\Theta(\frac{2nmp}{N})$ because of the used matrix partition approach shown in Figure 2. The I/O complexity of CO-matmul, however, is different from the basic algorithm: $\Theta(n+m+p+\frac{nm+mp+np}{B}+\frac{nmp}{B\sqrt[2]{Z}})$ [12]. Applying the ICE model to CO-matmul, the total energy complexity is computed as Equation 20.

Table 6: Matmul Complexity Analysis

Complexity	Cache-oblivious Algorithm	Basic Algorithm
Work	$\Theta(2nmp)$ [12]	$\Theta(2nmp)$ [33]
I/O	$\Theta(n+m+p+\frac{nm+mp+np}{B}+\frac{nmp}{B\sqrt[2]{Z}})$ [12]	$\Theta(\frac{nm+nmp+np}{B})$ [this study]
Span	$\Theta(\frac{2nmp}{N})$ [this study]	$\Theta(\frac{2nmp}{N})$ [this study]

6 Validation of ICE Model

This section describes the experimental study to validate the ICE model, including: introducing the two experimental platforms and how to obtain their parameters for the ICE model (cf. Section 6.1), describing SpMV implementation and sparse matrix types used in this validation (cf. Section 6.3), and discussing the validation results of SpMV and matmul.

6.1 Experiment Set-up

For the validation of the ICE model, we conduct the experiments on two HPC platforms: one platform with two Intel Xeon E5-2650l v3 processors and one platform with Xeon Phi 31S1P processor. The Intel Xeon platform has two processors Xeon E5-2650l v3 with 2×12 cores, each processor has the frequency 1.8 GHz. The Intel Xeon Phi platform has one processor Xeon Phi 31S1P with 57 cores and its frequency is 1.1 GHz. To measure energy consumption of the platforms, we read the PCM MSR counters for Intel Xeon and MIC power reader for Xeon Phi.

6.2 Identifying Platform Parameters

We apply the energy roofline approach [8, 7] to find the platform parameters for the two new experimental platforms, namely Intel Xeon E5-2650l v3 and Xeon Phi 31S1P. Moreover, the energy roofline study [7] has also provided a list of other platforms including CPU, GPU, embedded platforms with their parameters considered in the Roofline model. Thanks to authors Choi et al. [7], we extract the required values of ICE parameters for nine platforms presented in their study as follows: $\epsilon_{op} = \epsilon_d$, $\epsilon_{I/O} = \epsilon_{mem} \times B$, $\pi_{op} = \pi_1 \times \tau_d$, $\pi_{I/O} = \pi_1 \times \tau_{mem}$, where B is cache block size, ϵ_d, ϵ_d, τ_d, τ_{mem} are defined by [7] as energy per flop, energy per byte, time per flop and time per byte, respectively.

The ICE parameter values of the two new HPC platforms (i.e., Xeon and Xeon-Phi 31S1P) used to validate the ICE model are obtained by using the same approach as energy roofline study [8]. We create micro-benchmarks for the two platforms and measure their energy consumption and performance. The ICE parameter values of each platform are obtained from energy and performance data by regression techniques. Along with the two HPC platforms used in this validation, we provide parameters required in the ICE model for a total of 11 platforms. Their platform parameters are listed in Table 3 for further uses.

6.3 SpMV Implementation

We want to conduct complexity analysis and experimental study with two SpMV algorithms, namely CSB and CSC. Parallel CSB and sequential CSC implementations are available thanks to the study from Buluç et al. [6]. Since the optimization steps of available parallel SpMV kernels (e.g., pOSKI [1], LAMA[11]) might affect the work complexity of the algorithms, we decided to implement a

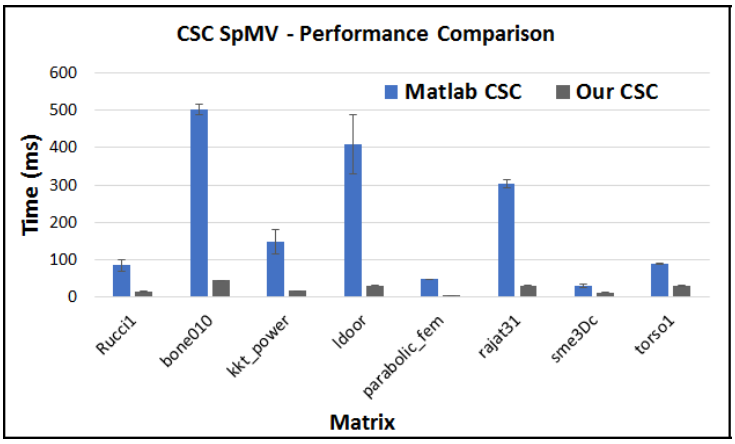

Figure 4: Performance (time) comparison of two parallel CSC SpMV implementations. For a set of different input matrices, the parallel CSC SpMV using Cilk out-performs Matlab parallel CSC.

simple parallel CSC using Cilk and pthread. To validate the correctness of our parallel CSC implementation, we compare the vector result y from $y = A * x$ of CSC and CSB implementation. The comparison shows the equality of the two vector results y. Moreover, we compare the performance of the our parallel CSC code with Matlab parallel CSC-SpMV kernel. Matlab also uses CSC layout as the format for their sparse matrix [14] and is used as baseline comparison for SpMV studies [6]. Our CSC implementation has out-performed Matlab parallel CSC kernel when computing the same targeted input matrices at least 136% across different inputs from Table 7. The experimental study of SpMV energy consumption is then conducted with CSB SpMV implementation from Buluç et al. [6] and our CSC SpMV parallel implementation.

6.4 SpMV Matrix Input Types

We conducted the experiments with nine different matrix-input types from Florida sparse matrix collection [10]. Each matrix input has different properties listed in Table 4, including size of the matrix $n \times m$, the maximum number of non-zero of the sparse matrix nz, the maximum number of non-zero elements in one column nc. Table 7 lists the matrix types used in this experimental validation with their properties.

6.5 Validating ICE Using Different SpMV Algorithms

The model aims to compare energy consumption of two algorithm. Therefore, we validate the ICE model by showing the comparison using the ratio of energy consumption of two algorithms. From the model-estimated data, CSB SpMV consumes less energy than CSC SpMV on both platforms. Even though CSB

Table 7: Sparse matrix input types. The maximum number of non-zero elements in a column nc is derived from [6].

Matrix type	**n**	**m**	**nz**	**nc**
bone010	986703	986703	47851783	63
kkt_power	2063494	2063494	12771361	90
ldoor	952203	952203	42493817	77
parabolic_fem	525825	525825	3674625	7
pds-100	156243	517577	1096002	7
rajat31	4690002	4690002	20316253	1200
Rucci1	1977885	109900	7791168	108
sme3Dc	42930	42930	3148656	405
torso1	116158	116158	8516500	1200

has higher work complexity than CSC, CSB SpMV has less I/O complexity than CSC SpMV. Firstly, the dynamic energy cost of one I/O is much greater than the energy cost of one operation (i.e., $\epsilon_{I/O} >> \epsilon_{op}$) on both platforms. Secondly, CSB has better parallelism than CSC, computed by $\frac{Work}{Span}$, which results in shorter execution time. Both reasons contribute to the less energy consumption of CSB SpMV. The measurement data confirms that CSB SpMV algorithm consumes less energy than CSC SpMV algorithm, shown by the energy consumption ratio between CSC-SpMV and CSB-SpMV greater than 1 in the Figure 5 and 6. For all input matrices, the ICE model has confirmed that CSB SpMV consumes less energy than CSC SpMV algorithm. Because the model has abstracted possible platform by only 4 parameters (i.e., ϵ_{op}, $\epsilon_{I/O}$, π_{op}, and $\pi_{I/O}$), there are the differences between the model and experiment ratios shown in the Figure 5 and 6. For accurate models that provide the precise energy estimation, the platform parameters need to be highly detailed such as RTHpower model for embedded platforms [28, 27].

6.6 Validating ICE With Matmul Algorithms

From the model-estimated data, Basic-Matmul consumes more energy than CO-Matmul on both platforms. Even though both algorithms have the same work and span complexity, Basic-Matmul has more I/O complexity than CO-Matmul, which results in greater energy consumption of Basic-Matmul compared to CO-Matmul algorithm. The measurement data confirms that Basic-Matmul algorithm consumes more energy than CO-Matmul algorithm, shown by the energy consumption ratio between Basic-Matmul and CO-Matmul greater than 1 in the Figure 7 and 8. For all input matrices, the ICE model has confirmed that Basic-Matmul consumes more energy than CO-Matmul algorithm.

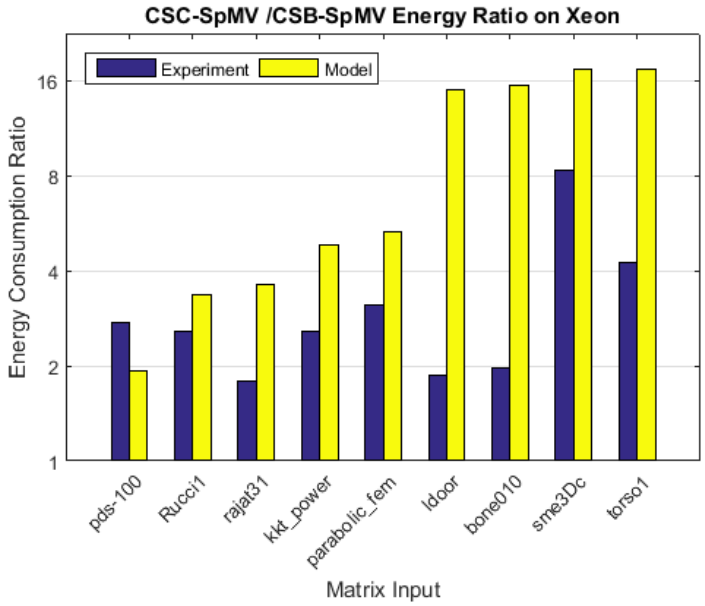

Figure 5: Energy consumption comparison between CSC-SpMV and CSB-SpMV on the Intel Xeon platform, computed by $\frac{E_{CSC}}{E_{CSB}}$. Both the ICE model estimation and experimental measurement on Intel Xeon platform show the consistent results that $\frac{E_{CSC}}{E_{CSB}}$ is greater than 1, meaning CSC SpMV algorithm consumes more energy than the CSB SpMV algorithm on different input matrices.

7 Related Work - Overview of energy models

Energy models for finding energy-optimized system configurations for a given application have been recently reported [12, 16, 19]. Imes et al. [16] used controller theory and linear programming to find energy-optimized configurations for an application with soft real-time constraints at runtime. Mishra et al. [24] used hierarchical Bayesian model in machine learning to find energy-optimized configurations. Snowdon et al. [26] developed a power management framework called Koala which models the energy consumption of the platform and monitors an application' energy behavior. Although the energy models for finding energy-optimized system configurations have resulted in energy saving in practice, they focus on characterizing system platforms rather than applications and therefore are not appropriate for analyzing the energy complexity of application algorithms.

Another direction of energy modeling study is to predict the energy consumption of applications by analyzing applications without actual execution on real platforms which we classify as analytic models.

Energy roofline models [8, 7] are some of the comprehensive energy models

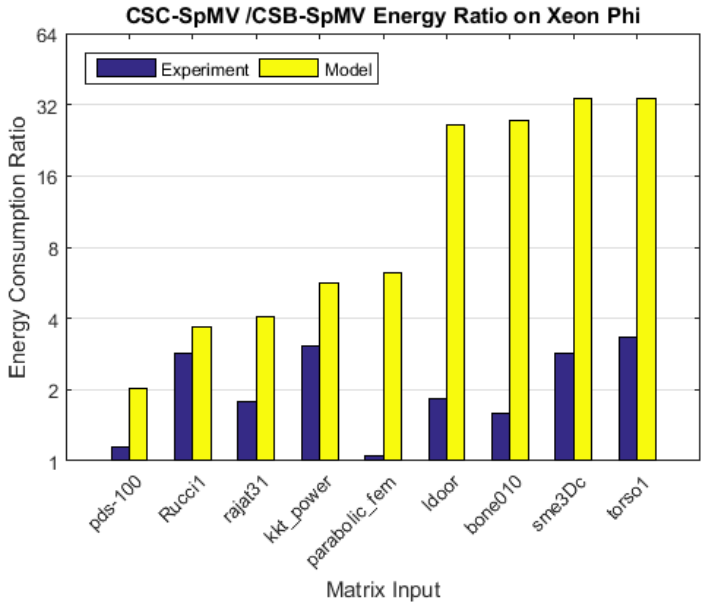

Figure 6: Energy consumption comparison between CSC-SpMV and CSB-SpMV on the Intel Xeon Phi platform, computed by $\frac{E_{CSC}}{E_{CSB}}$. Both the ICE model estimation and experimental measurement on Intel Xeon Phi platform show the consistent results that $\frac{E_{CSC}}{E_{CSB}}$ is greater than 1, meaning CSC SpMV algorithm consumes more energy than the CSB SpMV algorithm on different input matrices.

that abstract away possible algorithms in order to analyze and characterize different multicore platforms in terms of energy consumption. Our new energy model, which abstracts away possible multicore platform and characterize the energy complexity of algorithms based on their *work, span* and *I/O* complexity, complements the energy roofline models.

Validated energy models for *specific* algorithms have been reported recently [2, 23]. Alonso et al. [2] provided an accurate energy model for three key dense matrix factorizations. Malossi et al. [23] focused on basic linear-algebra kernels and characterized the kernels by the number of arithmetic operations, memory accesses, reduction and barrier steps. Although the energy models for specific algorithms are accurate for the target algorithms, they are not applicable for other algorithms and therefore cannot be used as general energy complexity models for parallel algorithms.

The *energy scalability* of a parallel algorithm has been investigated by Korthikanti et al. [17, 18]. Unlike the energy scalability studies that have not been validated on real platforms, our new energy complexity model is validated on HPC and accelerator platforms, confirming its usability and accuracy.

The energy complexity of *sequential* algorithms on a *uniprocessor* machine

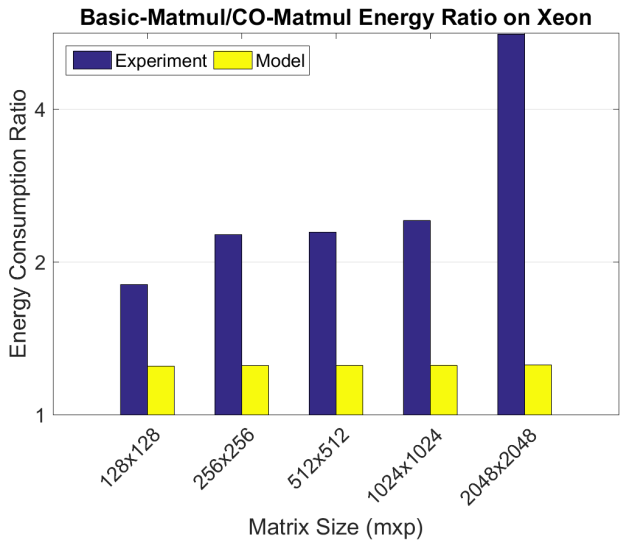

Figure 7: Energy consumption comparison between Basic-Matmul and CO-Matmul on the Intel Xeon platform, computed by $\frac{E_{Basic}}{E_{CO}}$. Both the ICE model estimation and experimental measurement on Intel Xeon platform show that $\frac{E_{Basic}}{E_{CO}}$ is greater than 1, meaning Basic-Matmul algorithm consumes more energy than the CO-Matmul algorithm.

with *several memory banks* has been studied by Roy et al. [25]. Our energy complexity studies complement Roy et al.'s studies by investigating the energy complexity of *parallel* algorithms on a *multiprocessor* machine with *a shared memory bank* and private caches, a machine model that has been widely adopted to study parallel algorithms [13, 3, 18].

8 Conclusion

In this study, we have devised a new general model for analyzing the energy complexity of multithreaded algorithms. The energy complexity of an algorithm is derived from its *work*, *span* and *I/O* complexity. Moreover, two case studies are conducted to demonstrate how to use the model to analyze the energy complexity of SpMV algorithms and matmul algorithms. The energy complexity analyses are validated for two SpMV algorithms and two matmul algorithm on two HPC platforms with different input matrices. The experimental results confirm the theoretical analysis with respect to which algorithm consumes more energy. The ICE energy complexity model gives algorithm-developers the insight into which algorithm is analytically more energy-efficient. Improving the ICE model by considering the numbers of platform cores is a part of our future work.

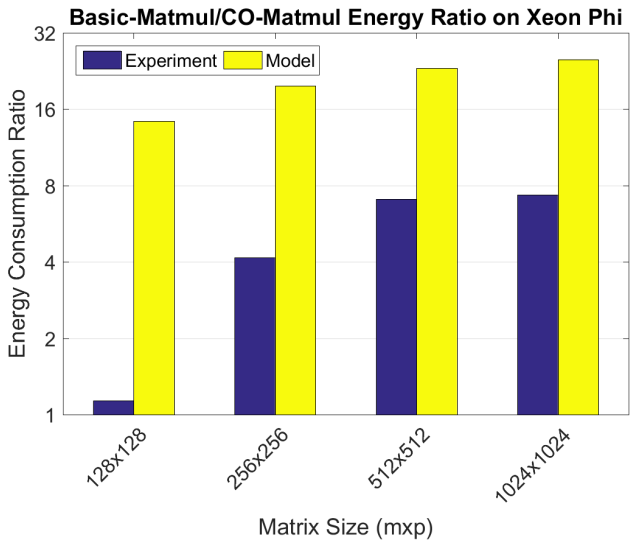

Figure 8: Energy consumption comparison between Basic-Matmul and CO-Matmul on the Intel Xeon Phi platform, computed by $\frac{E_{Basic}}{E_{CO}}$. Both the ICE model estimation and experimental measurement on Intel Xeon Phi platform show that $\frac{E_{Basic}}{E_{CO}}$ is greater than 1, meaning Basic-Matmul algorithm consumes more energy than the CO-Matmul algorithm.

References

[1] "Poski: Parallel optimized sparse kernel interface," http://bebop.cs.berkeley.edu/poski, accessed: 2015-11-17.

[2] P. Alonso, M. F. Dolz, R. Mayo, and E. S. Quintana-Orti, "Modeling power and energy consumption of dense matrix factorizations on multicore processors," *Concurrency Computat.*, 2014.

[3] L. Arge, M. T. Goodrich, M. Nelson, and N. Sitchinava, "Fundamental parallel algorithms for private-cache chip multiprocessors," in *Procs of the Twentieth Annual Symp on Parallelism in Algorithms and Architectures*, ser. SPAA '08, 2008, pp. 197–206.

[4] K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, and K. A. Yelick, "The landscape of parallel computing research: A view from berkeley," *Technical Report No. UCB/EECS-2006-183, University of California, Berkeley*, 2006.

[5] M. Bender, G. Brodal, R. Fagerberg, R. Jacob, and E. Vicari, "Optimal sparse matrix dense vector multiplication in the i/o model," *Theory of Computing Systems*, vol. 47, no. 4, pp. 934–962, 2010.

[6] A. Buluç, J. T. Fineman, M. Frigo, J. R. Gilbert, and C. E. Leiserson, "Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks," in *Procs of the Twenty-first Annual Symp on Parallelism in Algorithms and Architectures*, ser. SPAA '09, 2009.

[7] J. Choi, M. Dukhan, X. Liu, and R. Vuduc, "Algorithmic time, energy, and power on candidate hpc compute building blocks," in *Procs of the 2014 IEEE 28th Int Parallel and Distributed Processing Symp*, ser. IPDPS '14, 2014, pp. 447–457.

[8] J. W. Choi, D. Bedard, R. Fowler, and R. Vuduc, "A roofline model of energy," in *Procs of the 2013 IEEE 27th Int Symp on Parallel and Distributed Processing*, ser. IPDPS '13, 2013, pp. 661–672.

[9] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, Third Edition*. The MIT Press, 2009.

[10] T. A. Davis and Y. Hu, "The university of florida sparse matrix collection," *ACM Trans. Math. Softw.*, vol. 38, no. 1, pp. 1:1–1:25, 2011.

[11] M. Forster and J. Kraus, "Scalable parallel amg on ccnuma machines with openmp," *Computer Science - Research and Development*, vol. 26, no. 3-4, pp. 221–228, 2011. [Online]. Available: http://dx.doi.org/10.1007/s00450-011-0159-z

[12] M. Frigo, C. E. Leiserson, H. Prokop, and S. Ramachandran, "Cache-oblivious algorithms," in *Procs of the 40th Annual Symp on Foundations of Computer Science*, ser. FOCS, 1999.

[13] M. Frigo and V. Strumpen, "The cache complexity of multithreaded cache oblivious algorithms," in *Procs of the Eighteenth Annual ACM Symp on Parallelism in Algorithms and Architectures*, ser. SPAA '06, 2006, pp. 271–280.

[14] J. R. Gilbert, C. Moler, and R. Schreiber, "Sparse matrices in matlab: Design and implementation," *SIAM J. Matrix Anal. Appl.*, vol. 13, no. 1, pp. 333–356, Jan. 1992. [Online]. Available: http://dx.doi.org/10.1137/0613024

[15] P. Ha, V. Tran, I. Umar, P. Tsigas, A. Gidenstam, P. Renaud-Goud, I. Walulya, and A. Atalar, "Models for energy consumption of data structures and algorithms," EU FP7 project EXCESS deliverable D2.1 (http://www.excess-project.eu), Tech. Rep., 2014.

[16] C. Imes, D. Kim, M. Maggio, and H. Hoffmann, "Poet: a portable approach to minimizing energy under soft real-time constraints," in *Real-Time and Embedded Technology and Applications Symp (RTAS), 2015 IEEE*, 2015, pp. 75–86.

[17] V. Korthikanti and G. Agha, "Analysis of parallel algorithms for energy conservation in scalable multicore architectures," in *Int Conf on Parallel Processing, 2009. ICPP '09.*, 2009, pp. 212–219.

[18] V. A. Korthikanti and G. Agha, "Towards optimizing energy costs of algorithms for shared memory architectures," in *Procs of the Twenty-second Annual ACM Symp on Parallelism in Algorithms and Architectures*, ser. SPAA '10, 2010, pp. 157–165.

[19] V. Kotlyar, K. Pingali, and P. Stodghill, "A relational approach to the compilation of sparse matrix programs," Tech. Rep., 1997.

[20] P. Kumar, A. Gurtov, and P. H. Ha, "An efficient authentication model in smart grid networks," in *Procs of the 15th Int Conf on Information Processing in Sensor Networks*, 2016, pp. 65:1–65:2.

[21] J. Lagravire, J. Langguth, M. Sourouri, P. H. Ha, and X. Cai, "On the performance and energy efficiency of the pgas programming model on multicore architectures," in *Procs of the Int Workshop on Optimization of Energy Efficient HPC & Distributed Systems (OPTIM-HPCS)*, 2016, pp. 800–807.

[22] L. Li and C. Kessler, "Validating energy compositionality of GPU computations," in *HIPEAC Workshop on Energy Efficiency with Heterogeneous Computing (EEHCO)*, 2015.

[23] A. C. I. Malossi, Y. Ineichen, C. Bekas, A. Curioni, and E. S. Quintana-Orti, "Systematic derivation of time and power models for linear algebra kernels on multicore architectures," *Sustainable Computing: Informatics and Systems*, vol. 7, pp. 24 – 40, 2015.

[24] N. Mishra, H. Zhang, J. D. Lafferty, and H. Hoffmann, "A probabilistic graphical model-based approach for minimizing energy under performance constraints," in *Procs of the Twentieth Int Conf on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '15, 2015, pp. 267–281.

[25] S. Roy, A. Rudra, and A. Verma, "An energy complexity model for algorithms," in *Procs of the 4th Conf on Innovations in Theoretical Computer Science*, ser. ITCS '13, 2013, pp. 283–304.

[26] D. C. Snowdon, E. Le Sueur, S. M. Petters, and G. Heiser, "Koala: A platform for os-level power management," in *Procs of the 4th ACM European Conference on Computer Systems*, ser. EuroSys '09, 2009.

[27] V. N. N. Tran, B. Barry, and P. H. Ha, "Power models supporting energy-efficient co-design on ultra-low power embedded systems," in *Procs of the Intl Conf on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XVI), in press.*, 2016.

[28] ——, "Rthpower: Accurate fine-grained power models for predicting race-to-halt effect on ultra-low power embedded systems," in *Procs of the 17th IEEE Int Symp on Performance Analysis of Systems and Software*, 2016, pp. 155–156.

[29] I. Umar, O. J. Anshus, and P. H. Ha, "Deltatree: A locality-aware concurrent search tree," in *Procs of the 2015 ACM SIGMETRICS Int Conf on Measurement and Modeling of Computer Systems*, 2015, pp. 457–458.

[30] ——, "Effect of portable fine-grained locality on energy efficiency and performance in concurrent search trees," in *Procs of the 21st ACM SIGPLAN Symp on Principles and Practice of Parallel Programming*, 2016, pp. 36:1–36:2.

[31] ——, "Greenbst: An energy-efficient concurrent search tree," in *Proc. of the 22nd Intl. European Conf. on Parallel and Distributed Computing (Euro-Par 16)*, 2016, pp. 502–517.

[32] S. Williams, A. Waterman, and D. Patterson, "Roofline: An insightful visual performance model for multicore architectures," *Commun. ACM*, vol. 52, no. 4, pp. 65–76, 2009.

[33] K. Yelick. (2004, Sept) Cs 267 parallel matrix multiplication. [Online]. Available: http://www.cs.berkeley.edu/~yelick/cs267